

Advanced Software Fault Tolerance Strategies For Mission Critical Spacecraft Applications

for

NASA Ames Research Center/IV&V Facility
Software Initiative UPN 323-08

Task 2 Interim Report April 30, 1999
Common Errors Afflicting Real Time Control System Software

Prepared by JSC NX Technology Division

Approved by M. Himel/Chief, NX Technology Division

Purpose and scope

This interim report is an account of software anomalies encountered in real-time control systems. The majority of data collected is from the National Space Transportation System (Space Shuttle) program, but information was also collected for some “well-known” failures induced by software. While there is much anecdotal information available about the effects of software errors, which the interested reader should consult (Neumann 95, Wiener 93, Peterson 95), the data presented here is limited to that for which enough detail is provided so that a reasonable inference about the probable root cause could be drawn. Unfortunately, this rather restricts the number of anomalies available for analysis.

The software anomaly list for Space Shuttle missions was obtained from a special report generated by United Space Alliance at the request of the JSC Avionics Engineering Division. The anomaly list includes only those errors manifested in flight and only those anomalies that were the result of a design or code error. Anomalous behavior that was the result of a requirements deficiency are not “code breaks” by definition and so aren’t included.

The NASA Lessons Learned database was examined and very little useful information could be discovered. It appears that there is no systematic approach to collecting and distributing data relating to SW anomalies that extends throughout NASA. Given the large number of programs, the number of different software systems, and the different development approaches that NASA has developed through its history, such a central repository of software failure would be most useful in identifying weaknesses in software design, development, test, and evaluation activities.

For software programs developed outside of NASA, world wide web searches provided much information. Unfortunately most of it was of anecdotal character and is excluded from this report. Other sources that seemed promising, such as the NTSB, imposed very high search costs, i.e. one couldn’t search for software misbehavior as the probable cause and there were too many reports to read individually. The commercial world is even less open with accounts of software induced failures. It was only in highly publicized cases that such information could be found.

Findings

1. Incomplete or inconsistent requirements can impose inordinate costs on a program.
The high costs are imposed not only in the corrective actions required but also can be extended much further. Very simple errors in the requirements have, in more than one case, led to the loss of the vehicle as the most extreme consequence. Less spectacularly, but as serious, are the errors that have led to the loss of the mission.
2. Exhaustive testing fails to prevent errors in the operations environment.
The Space Shuttle’s on-board control system is one of the more exhaustively tested systems in use today. This testing regime has not eliminated all errors from the system. With systems becoming ever more complex, testing becomes an ever less efficient means of assuring correctness.

Listing of Software Induced Errors in Real-Time Control Systems

Table 1 is the definitive list of the Space Shuttle Discrepancy Reports (DRs) opened on anomalies that manifested themselves in-flight and were identified to be software induced. The listing was generated by United Space Alliance at the request of the JSC Avionics Engineering Division. The original list provided the DR number and title, the software system(s) in error, the severity of the consequences of the error, and the date the error was logged.

The logged date was not always during a flight. This could be the result of the way the Mission Operations Directorate (MOD) operates in flight mode: the anomaly would be logged at the time of its occurrence by MOD but not be logged as a DR until analysis had confirmed the particular software error responsible for the anomalous behavior.

The available paperwork supporting these DRs was analyzed and the probable root causes identified. For those cases where the paperwork was no longer available, the probable root cause was inferred from the title or, if this were not possible, left undetermined.

By definition, this list includes only those software errors that resulted from design or coding errors.

Table 2 lists several well known real-time control system failures, collected from a variety of sources, for which enough details were supplied that a probable root cause inference could be drawn.

Table 1 - National Space Transportation System SW Related In-flight Anomalies

Number	Severity	Title
DR028365		Potential Offset Between SIP Cycle & TSIP
DR047896		OMS to RCS Gauging Counter Error
DR051210		HUD Max Decel CMD
DR051347		Fail-to-Sync From DCI/DMC Contention for CDMB_MAT_CRT
DR054482		FDA Info Not Completely Deleted for L04P1202V
DR055055	3	Message Stacking Number Was Not Updated When Fault Summary Page Clear command is sent, counter goes to zero, variable is not reset until new error received.
DR055347	2	Orbiter P, Y, R, Rates Passed to Spacelab Incorrect Coded to input in RPY order instead of PYR
DR055355		Downlist I/O Cycle Wrap GPC Errors Cumulative GMT updates cause GPC errors
DR055798		SL Special Processing Not Being Called
DR055799		SL GNC Transfer GMT Days Incorrect
DR060552		Auto-MNVR Initialized By New Target Load
DR060587		Commanded Attitude Jump During LVLH – OEXDAP
DR061508		Limit Structures Incorrect for 2 S/L Parameters
DR062180	3	MFE Cycle Overruns in OPS 2 Mid-frequency executive cycle overruns occur during: Orbit operations, universal pointing in 3-axis rendezvous target track mode, both star trackers in star-of-opportunity star track mode, and data collection of a new star completed on phase 2 of the MFE. Load balancing reduced probability of reoccurrence
DR062967	3	Incorrect MPS Pitch Biases After a Pad Shutdown and a Hydraulic Failure At pad shutdown/abort, SSME pitch bias went to launch position, instead of remaining in gravity position. [Very similar to DR 101398]
DR100138		Incorrect Rel Nav Accumulated Velocities Not based solely on IMUs, not corrected for CG offset or filtered. Degradation of data, solution – user is not use rel nav display to monitor orbiter translational maneuvers
DR103036		Undesireable MPS CMD Configuration
DR103384	3	Unexpected “PBD CONFIG” Message During PBD Close

Advanced Software Fault Tolerance Strategies For Mission Critical Systems
Task 2 Interim Report - Common Errors Afflicting Real-time Systems

As payload bay doors are being closed, if commanded to stop, as doors stop, they engage ready-to-latch switches, then command to close, confused state

Table 1 - National Space Transportation System SW Related In-flight Anomalies (Continued)

Number	Severity	Title
DR103439	3	EIU I/O Errors And EIU Bypass Occurrence of errors is dependent on timing of Main Engine Controller Power Off WRT timing of I/O between the Main Engine Controller and EIU Logged during Ascent
DR103936	3	Unexpected PDRS Slip Annunciations During RMS Power-down Seems to be similar to PBD DR, when power switch set to off, transmission of position can be delayed by 1 cycle during which time the FSW continues to perform causing the data to become unreliable
DR104025		Incorrect Unit Vector From VV10D3 Erroneous SV uplink forced, prediction of SV to current time produced SV components 1-5 powers of 10 from overflowing.
DR194377	3	G9 FTS Due To PCMMU Errors PCMMU/GPC config gives each GPC its own bus to PCMMU, if I/O error, non-universal error, FTS caused by problems w/ PCMMU, not GPC
DR105363	1N	Manual Nose Wheel Steering (NWS) Activation For STS-60 NWS not activated automatically by FSW. Possible for FSW not to set ground speed enable flag for slow ground speed landings at nose gear slardown
DR107106	3	Incorrect MET Ref Time Displayed During pre-launch, MET is set ~ to GMT as a reference time calculated by subtracting MET from GMT. If MET > GMT, negative values occur. MET is reset at SRB ignition. Not seen by crew or KSC, only by MCC
DR107737	3	Init T/O Value Too Long For GTG Overlay If HW failure precludes successful I/O transaction on MMU bus, master seq controller T/O error logged instead of ITO if initial iteration count for GTG overlays is < ITO value. GTG overlay completes successfully on alt MMU bus, provided no failures there.
DR107772	3	Transient IMU Resolver Limit BITE During STS-63 Ground Checkout - IMU HW executed high rate gimbal flips causing extremely rapid resolver movement such that gimbal angle rates could not be accurately determined by FSW leading to transient IMU resolver limit BITE
DR107971	3	False OFF/BSY MMU Message Possible - If Mass Memory transaction requested w/n 160 msec of last completed transaction, may be rejected (OFF/BUSY MMUX) due to FCOS incorrectly indicating busy.
DR109204	3	Possible Negative Duty Cycle During GPC Initialization

Table 1 - National Space Transportation System SW Related In-flight Anomalies (C

Number	Severity	Title
DR109209	3	Suspended Attitude Calculation Interferes With Antenna
DR110271		Incorrect Initialization of Number of Jets For PTI PASS FCS parameters for # of yaw RCS jets used by WRAP DAP are initialized to default values rather than appropriate I-LOADS
DR110288	2	Incorrect Switch Conditions For DNYP Roll Rate Lat Accel Component should be computed by the DNYP_COMP function when either LATE or PTI_WRAP flag is true. Existing FSW only checks for LATE flag, so if PTI_WRAP flag is set, DNYP_COMP is not used for computation.
DR110472	2	Cyclic GPC Errors Can Block GNC-To-SM SV ICC Transfer GPC errors were occurring at 6.25 Hz, leaving 40 HW (word-length) available for the SV to be transferred. The SV is 42 HW in length, so it was continually prevented from being transferred until the GPC errors were stopped.

NOTES:

- STS-xx? denotes that the flight was not identified in either the DR or the supporting paperwork and represents the flight nearest to the error.
- Severity Codes are:
 1. Loss of vehicle / Loss of crew
 2. Loss of Mission
 3. Significant Mission Impact - Workaround available
 4. Insignificant Impact - Workaround available
 5. No impact - Paperwork, standards violations, etc
- Severity 1N (DR105363) means that the consequences of the software error are Loss of vehicle / Loss of crew but that established procedures prevent those consequences from ever occurring - in this case, the nosewheel steering was activated manually.

Table 2 - Other “well known” real-time control system software anomalies:

Description	Date	Probable Root Cause
Apollo 11 LM <i>Eagle</i> manifested several “1201” (overflow) alarms during powered descent. No HW input sanity check: the rendezvous radar was seeking an angle that had a sine and a cosine of 0. (Neumann 95, Wiener 93)	20 July 1969	Design Error
Therac-25: a linear accelerator used for cancer treatments overdosed six people. (Leveson 95)	June 1985 - January 1987	Incomplete/ Incorrect Requirements; Design/code errors
AT&T Network Outage (Neumann 95)	15 January 1990	Code error
STS-49 on-board software failed to calculate maneuver parameters for <i>Intelsat 6</i> rendezvous because of a precision mismatch in the Newton-Raphson iteration (Neumann 95 as well as personal experience)	May 1992	Incomplete/ Incorrect Requirements
Ariane V loss - IMU alignment routines continued after lift-off; arithmetic overflow sent box to diagnostic; FC computer interprets diagnostics as valid measurements (ARIANE 5, Flight 501 Failure , Report by the Inquiry Board, 19 July 1996 , Paris)	4 June 1996	Incomplete/ Incorrect Requirements
Mars Pathfinder: loss of data/continual restarts. Priority inversion led to deadlock; watchdog restarted system. (RISKS 19:49; 19:50; 19:53)	July 1997	Design/code error
Lewis Spacecraft Loss: thruster imbalance initiated spin in X axis, which was undetected by the two-axis gyro system, and coupled to Z axis. Attitude control system shutdown after “too many firings” to control the Z axis spin , which caused the solar arrays to be unilluminated, led to complete discharge of the batteries. (Lewis Spacecraft Mission Failure Investigation Board Final Report, 12 February 1998)	26 August 1997	Incomplete/ Incorrect Requirements
Titan IVA Loss - FC computer re-initialize after interruption of electrical power; initial state did not match external state. (Titan 4A ,Air Force Space Command News Release, Jan. 15, 1999)	12 August 1998	Incomplete/ Incorrect Requirements
Delta III Loss - 4Hz torsional mode missing from flight control model; simulation had shown the mode was not “significant.” (Delta III, Boeing Press Release, October 19, 1998, as well as personal communication with failure board members)	26 August 1998	Incomplete/ Incorrect Requirements

Extracted from the NASA Lessons Learned Database

Lesson Number: 0288

Subject/Title/Topic(s): Galileo Spacecraft Safing During Star Scanner Calibration

Description of Driving Event :

An Attitude and Articulation Control Subsystem (AACS) sequence designed to collect data for calibration of the spacecraft star scanners in the AACS inertial mode (gyros on), was tested on the Galileo test bed simulator.

Prior to the transmission and execution of this calibration sequence on the spacecraft, a star misidentification event caused the AACS to switch from the "inertial" to the "cruise" mode (gyros off). Because of the importance of getting the calibration data and limited open time in the few weeks before Venus encounter, it was decided to proceed with the star scanner calibration. The mode change was evaluated and not believed to have an effect on the planned sequence.

However, during the execution of the calibration sequence, a spin bearing controller instability occurred due to an unexpected incompatibility between the mode and AACS software. This caused a series of hardware swaps within the AACS, ultimately causing the spacecraft to go into safing. A subsequent test on the Galileo test bed simulator duplicated the spacecraft response.

This event occurred twenty-five days before Venus encounter and the difficult recovery process from safing took three weeks. Had this anomaly occurred closer to the encounter, significant impact on science data return could have resulted.

Lesson Number: 0310

Subject/Title/Topic(s): Mars Observer Inertial Reference Loss

Description of Driving Event:

Mars Observer experienced inertial reference loss on several occasions during its cruise to Mars. Two classes of inertial reference loss have been observed:

A. In early January 1993, the flight software was unable to identify any star that transited the celestial sensor assembly field of view. The unidentified stars count exceeded the "loss logic limit," and the fault protection software commanded the spacecraft to the sun coning attitude contingency mode. This occurred three times before a temporary software script to widen the star identification tolerance was uplinked in order to artificially increase the attitude uncertainties, or covariances, used by the software. Design flexibility of the flight computer and software allowed the software patch to be easily performed. It was suspected that the cause was due to the use of the more optimistic gyro noise parameters and values obtained from the in-house test results rather than the manufacturer's specifications. Recovery time: 3 days per occurrence.

B. During April and May 1993, three more incidents caused the spacecraft to declare inertial reference loss when the "sun monitor ephemeris" test, which compares the expected new position with the measured positions, was violated. An algorithm error in the inherited flight software caused the spacecraft attitude to be incorrectly estimated under certain conditions. A similar problem occurred on the Defense Meteorological Satellite Program (DMSP), an earth orbiting spacecraft built by the same contractor, that was using the same flight software. This algorithm error puts the spacecraft in additional jeopardy when the attitude covariances are large. Since the script that was intended to prevent the January incidents increases the covariances, the script acted as a catalyst for the three April/May anomalies. The review of the data indicated that no detailed code walk-through was performed on the software patch that widened the star identification tolerance. Recovery time: 5 days per occurrence.

Lesson Number: 0391

Subject/Title/Topic(s): Galileo Spacecraft Safing Recovery Anomalies

Description of Driving Event:

Two Galileo Spacecraft anomalies occurred during the week of September 12, 1994. Several difficulties were encountered during the recovery from these anomalies.

The first anomaly was a Data Bulk Unit Memory (DBUM) parity error on the Galileo Command and Data Subsystem (CDS). This nonprivileged error resulted in spacecraft safing but did not bring down either CDS string. In response to this anomaly, the flight team developed a special privileged command program to isolate the failed memory byte. Although this program was almost identical to a recent successfully run program, it required a nonstandard ground system configuration for command translation. This configuration was not established, causing the second anomaly which brought down the CDS A-string and re-executed safing.

During the recovery from this second safing, an existing recovery file was left unchanged from a year earlier, even though the one-way light time had doubled. As a result, an inappropriate Telemetry Modulation Unit (TMU) modulation index command was uplinked, causing a short data outage.

During both recovery efforts, it was believed that the system fault protection associated with safing had turned off the High Voltage (HV) to the Heavy Ion Counter (HIC). In fact, a patch had been implemented a year earlier which prevented the fault protection from turning off the HIC HV. A lack of proper configuration management along with an inadequate check of the spacecraft state created this confusion. Since a prolonged lack of HV to the HIC would permanently damage the HIC, the project was forced to consider a risky proposal to switch the HIC to the only remaining working CDS string in an unnecessary attempt to get the HIC HV back on.

The overall recovery time was 12 days. Had the spacecraft been in a critical operation mode, this long recovery time could have been more detrimental. A standardized anomaly recovery plan would have helped avoid some of the above problems.

Lesson Number: 0405

Subject/Title/Topic(s): Stepping Commands Cause Positioning Problems (1970/76)

Description of Driving Event :

Mariner (M)'69 and Viking Orbiter (VO)'75 had problems because the designs used stepping commands rather than explicit commands. Explicit commands cause a device (such as the scan platform) to move directly to the desired state or position, regardless of its current position. Stepping or incremental commands, however, simply cause the device to move by the commanded increment from its current position to the new position. If the current position is incorrect, the next position and all subsequent positions will be incorrect.

When a power problem occurred as M'69 was approaching its Mars fly-by, the scan platform moved inadvertently and the engineering telemetry needed to determine the platform's position was destroyed. The platform then had to be commanded by trial-and-error until Mars appeared in the returning pictures. Had this search not been successful in the short time remaining until Mars encounter, the entire science sequence would have been irretrievably lost.

A less serious problem occurred on VO-1 when the position of the camera filter wheel got out of step with the picture-taking sequence. Once out of step, the several frames subsequently exposed were lost due to over/under-exposure. Correct exposure had to be restored by ground command.

Analysis of Identified Software Induced Errors in Real-Time Control Systems

The limited amount of the data collected prevents any detailed analysis. The shuttle data is limited to design/code errors that were not detected in testing; analogous data from requirements failures is not available. The other real-time system failures are known because of their public or spectacular nature. This limits the available data to that of a particular failure; the other failures that the systems may have experienced are not available. Because of this bias, no valid analysis is possible.

References

- Leveson 95 - (Leveson, N., *Safeware: System Safety and Computers*, Addison-Wesley, 1995)
Neumann 95 - Neumann, P. G., *Computer Related Risks*, Addison-Wesley, 1995
Peterson 95 - Peterson, I., *Fatal Defect, Chasing Killer Computer Bugs*, Times Books, 1995
RISKS - Risk Forum Archive (archive of Usenet newsgroup comp.risks moderated by P. G. Neumann)
available at <http://catless.ncl.ac.uk/Risks/>
Wiener 93 - Wiener, L. R., *Digital Woes, Why We Should Not Depend on Software*, Addison-Wesley,
1993
NASA Lessons Learned Database - <http://llis.gsfc.nasa.gov/>
Orbital Anomalies in Goddard Spacecraft (OAGS) Annual Reports:
<http://arioch.gsfc.nasa.gov/302/oags.htm>